Neuro-GSTH: A Geometric Scattering and Persistent Homology Framework for Uncovering Spatiotemporal Signatures in Neural Activity

Dhananjay Bhaskar,[1†] **Jessica Moore**,[2†] **Yanlei Zhang**,[3] **Feng Gao**,[4] **Bastian Rieck**,[5] **Guy Wolf**,[3] **Helen Pushkarskaya**,[6] **Firas Khasawneh**,[7] **Elizabeth Munch**,[8-9] **Valentina Greco**,[1,10-13*] **Christopher Pittenger**,[1,10-13*] **Smita Krishnaswamy**[1,13-17*]

[1] Department of Genetics, Yale School of Medicine; [2] Department of Cell Biology, Duke University; [3] Mila - Quebec AI Institute; [4] Mailman School of Public Health, Columbia University; [5] Institute of AI for Health, Helmholtz Munich, Technical University of Munich; [6] Department of Psychiatry, Yale School of Medicine; [7] Department of Mechanical Engineering, Michigan State University; [8] Department of Computational Mathematics, Science and Engineering, Michigan State University; [9] Department of Mathematics, Michigan State University; [10] Department of Cell Biology, Yale School of Medicine; [11] Department of Dermatology, Yale School of Medicine; [12] Yale Stem Cell Center, Yale University; [13] Yale Cancer Center, Yale University; [14] Department of Computer Science, Yale University; [15] Program for Computational Biology and Bioinformatics, Yale University; [16] Applied Mathematics Program, Yale University; [17] Wu Tsai Institute, Yale University.

Correspondence: `smita.krishnaswamy@yale.edu`

[†] Co-first authors

Abstract

Understanding how neurons communicate and coordinate their activity is essential for unraveling the brain's complex functionality. To analyze the intricate spatiotemporal dynamics of neural signaling, we developed *Geometric Scattering Trajectory Homology (neuro-GSTH)*, a novel framework that captures time-evolving neural signals and encodes them into low-dimensional representations. GSTH integrates geometric scattering transforms, which extract multiscale features from brain signals modeled on anatomical graphs, with t-PHATE, a manifold learning method that maps the temporal evolution of neural activity. Topological descriptors from computational homology are then applied to characterize the global structure of these neural trajectories, enabling the quantification and differentiation of spatiotemporal brain dynamics.

We demonstrate the power of neuro-GSTH in neuroscience by applying it to both simulated and biological neural datasets. First, we used neuro-GSTH to analyze neural oscillatory

behavior in the Kuramoto model, revealing its capacity to track the synchronization of neural circuits as coupling strength increases. Next, we applied neuro-GSTH to neural recordings from the visual cortex of mice, where it accurately reconstructed visual stimulus patterns such as sinusoidal gratings. Neuro-GSTH-derived neural trajectories enabled precise classification of stimulus properties like spatial frequency and orientation, significantly outperforming traditional methods in capturing the underlying neural dynamics. These findings demonstrate that neuro-GSTH effectively identifies neural motifs—distinct patterns of spatiotemporal activity—providing a powerful tool for decoding brain activity across diverse tasks, sensory inputs, and neurological disorders. Neuro-GSTH thus offers new insights into neural communication and dynamics, advancing our ability to map and understand complex brain functions.

1 Introduction

The intricate architecture of the brain's neuronal circuits is fundamental to understanding its complex functioning. At the heart of this complexity lies neuronal signaling, a process that encompasses a range of dynamic patterns, including self-loops and oscillatory behavior. These patterns are not just prevalent in active brain states but also evident in the resting state of the brain, illustrating the continuous interplay of signals even in the absence of external stimuli. Moreover, stimulus-evoked activity adds another layer of complexity, showcasing how external inputs can modulate these intrinsic signaling dynamics. The investigation into these processes is crucial for unraveling the mysteries of brain functionality and disorders.

However, a significant challenge arises when attempting to quantify and visualize the high-dimensional spatiotemporal data obtained from brain imaging techniques. The sheer volume and complexity of this data pose a formidable obstacle for signal processing. Traditional methods are unable to capture intricate dynamics generated through interactions between neurons at the individual, circuit and systems scales. This difficulty is not just a technical barrier but also a bottleneck in advancing data-driven modeling in neuroscience. Effective visualization and quantification of this data are pivotal in developing models that can predict, classify, and simulate brain activities under various conditions, thereby opening new avenues in both research and clinical applications.

Recently, geometric and topological methods have emerged as promising approaches for quantifying the shape and structure of data at multiple scales of granularity. These approaches are particularly suited to tackle the complexity of brain signals, as they can effectively capture the intricacies of signaling patterns through multiple spatial and temporal scales. Motivated by applications of geometric scattering and persistent homology in the analysis of timelapse microscopy data, we sought to develop a technique that can be used to obtain low-dimensional representations of brain activity patterns. Such representations are useful for various downstream tasks, such as the identification of neural signaling motifs associated with the processing of visual stimuli, episodic memory retrieval, and the diagnosis of mental disorders.

Our method, neuro-GSTH (Geometric Scattering Trajectory Homology), employs a four-

step method to capture the spatiotemporal dynamics of neural activity. First we construct a spatial graph from the placement of sensors on the brain, with the time-lapse signal as node features. We then apply geometric scattering to obtain a multi-resolution descriptor of the spatial patterning across the brain regions. Geometric scattering employs wavelet-based transformations to extract multiscale representations of the signaling data, capturing the intricate hierarchical structures present in the brain structure. Next, we obtain a low dimensional trajectory of the signaling dynamics using TPHATE, a manifold-geometry preserving dimensionality reduction technique. The trajectory is generated by embedding geometric scattering features for each time point. Finally, we compute path signatures to obtain a multi-scale representation of the temporal dynamics in the TPHATE trajectory. Path signatures transform a path into a series of statistics or features, capturing various aspects of its geometry, such as total displacement, area enclosed, and loopiness. Path signatures are adept at dealing with irregularly sampled data, and are invariant to certain transformations of the path, such as reparametrization, making them robust to changes in the speed of traversal along the path.

We evaluated neuro-GSTH using a variety of mathematical models, simulations, and experimental data. First, we demonstrated that neuro-GSTH representations could predict the parameters of the Kuramoto model, simulating oscillatory behavior generated by excitatory-inhibitory coupling between neurons. Next, we applied neuro-GSTH to Ca^{2+} imaging data from layer 4 of the primary visual cortex in mice, showing that it can accurately reconstruct spatial frequency patterns of sinusoidal visual stimuli. Lastly, we used neuro-GSTH to analyze resting-state brain dynamics before and after a risky decision-making task, distinguishing between typically developed subjects and those with obsessive-compulsive disorder (OCD). Our results revealed that neuro-GSTH's attention mechanism focused on regions active in the default mode network (DMN), with persistent changes in attention scores after the task, indicating alterations in DMN activity.

Together, these results demonstrate the broad applicability of neuro-GSTH in analyzing complex spatiotemporal neural activity across a range of conditions and subject cohorts. By effectively capturing both local and global dynamics of brain activity, neuro-GSTH provides a powerful framework for identifying distinct neural activity motifs associated with cognitive tasks, sensory stimuli, and neurological disorders. The ability to model and classify brain activity patterns at multiple scales opens new avenues for understanding the mechanisms underlying neurodevelopmental conditions like autism spectrum disorder (ASD), as well as other brain disorders. Furthermore, the framework's capacity to fit mathematical models to neural signaling offers potential for developing predictive tools that can inform interventions and therapeutic strategies. Overall, neuro-GSTH lays the groundwork for integrating topological and geometric methods into neuroscience, paving the way for more robust and interpretable analyses of brain dynamics in both research and clinical applications.

2 Results

2.1 GSTH architecture and overview

Neuro-GSTH is an extension of GSTH, which we previously used to quantify Ca^{2+} signaling patterns in stem cells of the mouse epidermis (Moore et al., 2021). GSTH consists of three steps: (1) spatial embedding of the signal using geometric scattering, (2) temporal embedding of the multiscale geometric features using tPHATE to obtain a low-dimensional trajectory, and (3) computation of topological features of the trajectory using persistent homology. These steps are detailed in the following sections.

2.1.1 Learnable Geometric Scattering Module

First, we encode spatial variations in the signaling pattern at a given timepoint, t_i, using geometric scattering, an unsupervised method for generating embeddings for graph-structured data. This involves multiscale diffusion of the signal, $X(t_i)$, by performing random walks on a graph, $G(V, E)$, constructed from the spatial position of individual neurons or sensors in the input data. The graph consists of vertices, $V = \{v_1, v_2, \cdots, v_M\}$ and edges connecting vertices that lie within some predefined neighborhood distance. We represent the graph using the adjacency matrix, A, and the degree matrix, D. We use the notation $x(v_l, t_i)$ to denote the signaling activity at vertex v_l at time t_i. The design of graph wavelets is based on the diffusion operator constructed from lazy random walks, $R = \frac{1}{2}\left(I + AD^{-1}\right)$, over a graph, i.e.,

$$\Psi_0 = I - R\,, \quad \Psi_j = R^{2^{j-1}} - R^{2^j} = R^{2^{j-1}}(I - R^{2^{j-1}})\,, \quad j \geq 1 \tag{1}$$

The multi-scale nature of graph wavelets allows the geometric scattering module to traverse the entire graph in one layer, thus encoding both local and global graph features. Summation of the signal responses is used to obtain invariant graph-level features. Since the summation operation suppresses high frequency information, we complement it by using higher order summary statistics of signal x. Due to the iteration of applying graph wavelets followed by activation functions, the scattering module can be constructed as in a multi-layer (or multi-order) architecture. Specifically, the zeroth-order coefficients are calculated by taking statistical moments of the summation of signals, and the first-order features are obtained by applying a graph wavelet, which aggregates multiscale information of the graph. Second-order scattering features can further augment first-order features by iterating the graph wavelet and activation functions. The collection of scattering features provides a rich set of multiscale invariants of the graph G and can be used under both supervised and unsupervised settings for graph embedding.

For a signal $X(t_i) = [x(v_1, t_i), x(v_2, t_i), \ldots, x(v_M, t_i)]$ we compute the zeroth-order coefficients for each vertex/cell for timepoint t_i as follows:

$$S_0(X(t_i)) = \{R^{2^J} x(v_\ell, t_i) \quad \forall v_\ell \in V\} \tag{2}$$

The diffusion operator (lazy random walks) R works as a low pass filter that performs local averaging of signal in neighboring cells. Unlike the summation operator that averages

all vertex information and suppresses the high frequency information and hence has to be retrieved by higher order statistical moments, this retains finer description of cell/vertex embeddings. Then, by concatenating the wavelet coefficients for each cell/vertex at timepoint t_i, we can obtain the corresponding timepoint embedding $S_0(X(t_i))$ for timepoint t_i. Finally, the timepoint embedding for N timepoints can be calculated and the resulting $S_0(X(t)) = \{S_0(X(t_0)), S_0(X(t_1)), \ldots, S_0(X(t_n))\}$ is a feature matrix of dimension $N \times M$, where N is the number of timepoints and M is the number of vertices in graph. We hence obtain the zeroth-order coefficients for the N timepoints.

The zeroth-order features can be augmented by first-order features by applying graph wavelets and extracting finer description of high frequency response of a signal $X(t_i)$. Specifically, the first-order coefficients for each time point at each vertex are calculated as:

$$S_1(X(t_i)) = \{R^{2^J}\sigma_j x(v_\ell, t_i) \quad \forall v_\ell \in V, 1 \leq j \leq J\} \tag{3}$$

The value $\Psi_j x(v_\ell, t_i)$ aggregates the signal information $x(v_m, t_i)$ from the vertices v_m that are within 2^j steps of v_ℓ. It responds to sharp transitions or oscillations of the signal x within the neighborhood of v_ℓ with radius 2^j (in terms of the graph path distance). By concatenating all the vertex/cell embeddings, we can obtain the first order coefficients $S_1(X(t_i))$ for timepoint t_i.

Finally, the second-order coefficients can be obtained by further applying graph wavelets and extract even finer description of high frequency response of the signal $X(t_i)$:

$$S_2(X(t_i)) = \{R^{2^J}\sigma_{j'}\sigma_j x(v_\ell, t_i) \quad \forall v_\ell \in V, 1 \leq j < j' \leq J\} \tag{4}$$

These calculations are performed independently for each timepoint, resulting in N outputs. The first-order and second-order scattering coefficients generate a feature matrix of shape $N \times (M \times J)$ and $N \times (M \times 1/2(J \times (J-1)))$, respectively, as timepoint embeddings for the N timepoints. Finally, the zeroth-order, first-order and second-order coefficients were combined together as the embeddings for each time point $S(\mathbf{x}(t_i))$.

The geometric scattering module applies a cascade of graph wavelet transforms followed by an activation function. We typically use the absolute value function, $|\dot{|}$, although any another activation function can also be used.

2.1.2 tPHATE trajectories

The timelapse signal embeddings $S(X(t_i))$ from geometric scattering form a matrix of dimensions $T \times M$, where T is the number of time points in the data and M is the number of scattering coefficients for each time point. To visualize these embeddings, we calculate a distance matrix $D(i,j) = \|S(X(t_i)) - S(X(t_j)\|_2$ using the Euclidean distance between time point embeddings. D is converted from a distance matrix into a local affinity matrix K using an adaptive bandwidth Gaussian kernel, to capture local neighborhoods in the data:

$$K(x,y) = \frac{G(x,y)}{\|G(x,\cdot)\|_1^\alpha \|G(y,\cdot)\|_1^\alpha}, G(x,y) = e^{-\frac{\|x-y\|^\alpha}{\sigma}} \tag{5}$$

K is then row-normalized to define transition probabilities into the row stochastic matrix, P:

$$P(x, y) = \frac{K(x, y)}{\|K(x, \cdot)\|_1} \tag{6}$$

The probabilities P are then used for the Markovian random-walk diffusion process. The diffusion timescale, t_D, is then computed, which specifies the number of steps taken in the random-walk process. This parameter provides a tradeoff between encoding local and global information in the embedding, where a larger t_D corresponds to more steps than a smaller t_D. t_D is computed automatically using the spectral or von Neumann entropy of the diffusion operator. Next, we compute the diffusion potential distance between the distributions in the i-th and j-th rows of P^{t_D}, the t_D-step random walk over P:

$$P_D(i, j) = \sqrt{\sum_k (\log(P^{t_D}(i, k) - P^{t_D}(j, k)))^2} \tag{7}$$

Next, we compute the autocovariance of the $T \times M$ timelapse signal embeddings using $T-1$ lags, resulting in a $M \times (T-1)$ matrix. Next we average across the scattering coefficients M to obtain a single vector c of autocorrelation at each lag. We then construct the temporal affinity matrix A, by identifying the lag l_{max} where $c = 0$:

$$A(i, j) = \begin{cases} c(i, j), \text{ if } 0 < |i - j| < l_{\text{max}} \\ 0, \text{ otherwise} \end{cases} \tag{8}$$

We convert the autocorrelation matrix into the transition probability matrix P_T by row-normalizing the affinity matrix A and powering it to t^D. Finally, we alternate the spatial and temporal diffusion operators to construct a dual-view diffusion operator $P = P_D P_T$. We embed P using with metric MDS to obtain a 3-D tPHATE coordinates, $E(t) = (E_1(X(t)), E_2(X(t)), E_3(X(t)))$ for each time point t for visualization, which can also be viewed as point cloud $E = \{E(t_1), E(t_2), \ldots E(t_n)\}$ for subsequent topological data analysis.

2.1.3 Persistent homology of tPHATE trajectories

To obtain a topologically invariant characterization of a given tPHATE trajectory, denoted E, we calculated its persistent homology. Specifically, we calculated the persistent homology of E via a Vietoris–Rips filtration $\text{VR}_s(E)$. The Vietoris–Rips complex of E is defined as the filtered simplicial complex that contains a subset of E as a simplex if and only if all pairwise distances in the subset are less than or equal to s, i.e., $\text{VR}_s(E) = \{\{n_0, ..., n_m\} \mid \forall i, j \; d(i, j) \leq s\}$.

As described above, from $\text{VR}_s(E)$, we obtain a set of persistence diagrams Q consisting of birth-death-dimension triples $[b, d, q]$ that describe multiscale topological features of E. Each such point corresponds to a topological feature in the trajectory, which appears at a certain birth time and disappears at a death time. Note that the times are supposed to be understood with respect to the parameter s from above. The persistence of a point, $d - b$, therefore represents the

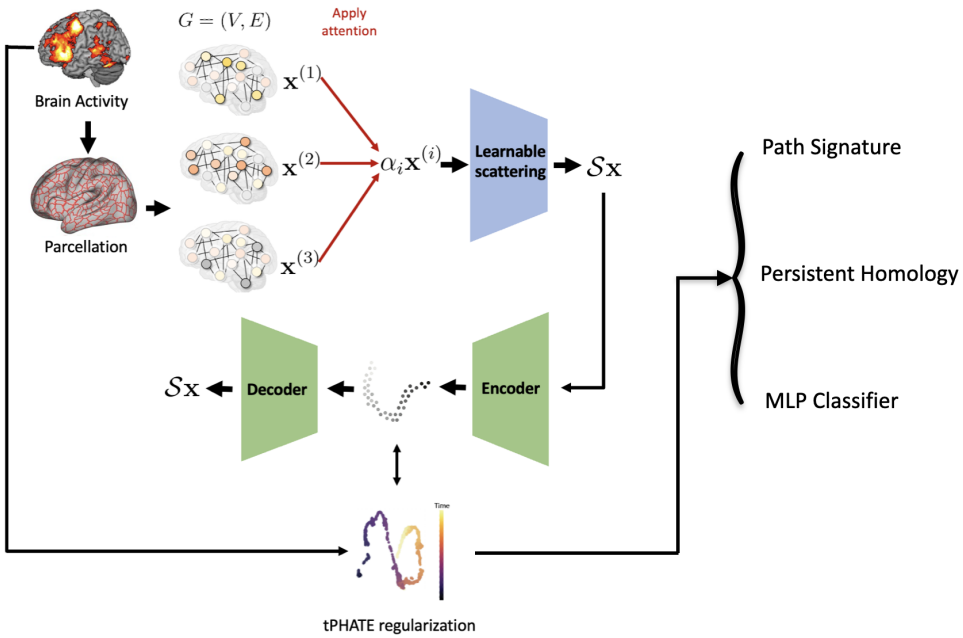

Figure 1: Neuro-GSTH architecture

prominence of the associated topological feature; higher values indicate that the feature occurs over a large scale, thus increasing its significance. We further calculated the associated Betti curves for each Q, resulting in a simple summary curve $B(Q, q)$ for the qth dimension consisting of the number of points (b_i, d_i) in Q such that $b_i \leq s < d_i$. The Betti curve characterizes the connectivity of $\text{VR}_s(E)$.

2.2 Comparisons and ablations

2.2.1 Takens embedding and time-varying topological signatures

While kymographs can be useful summaries of spatiotemporal patterns in data, they are typically designed in a problem-specific manner and with some knowledge of the expected dynamics. A more agnostic and less application-specific is to quantify the shape of the underlying data manifold using topological data analysis. Specifically, we encode the persistent homology of cell signals by constructing a point cloud using Taken's time-delay embeddings, and transforming the point cloud into a Vietoris-Rips simplicial complex.

We encode the signaling activity at each cell using a 2D time-delay positional embedding, $(X(v_l, t_i - \delta), X(v_l, t_i))$. Given a population of cells with signaling activity changing over time, we then construct a Betti curve from the 2D point cloud of the time-delay embeddings of cells at each timepoint t_i, and concatenate these into a matrix. We discretize time, t, along the columns and the filtration parameter, ϵ, along the rows. Entries in the matrix are the kth dimensional Betti number, β_k, for specific values of (t, ϵ). The matrix is a topological signature of the time-varying data and it can be vectorized for input into machine learning algorithms. The matrix

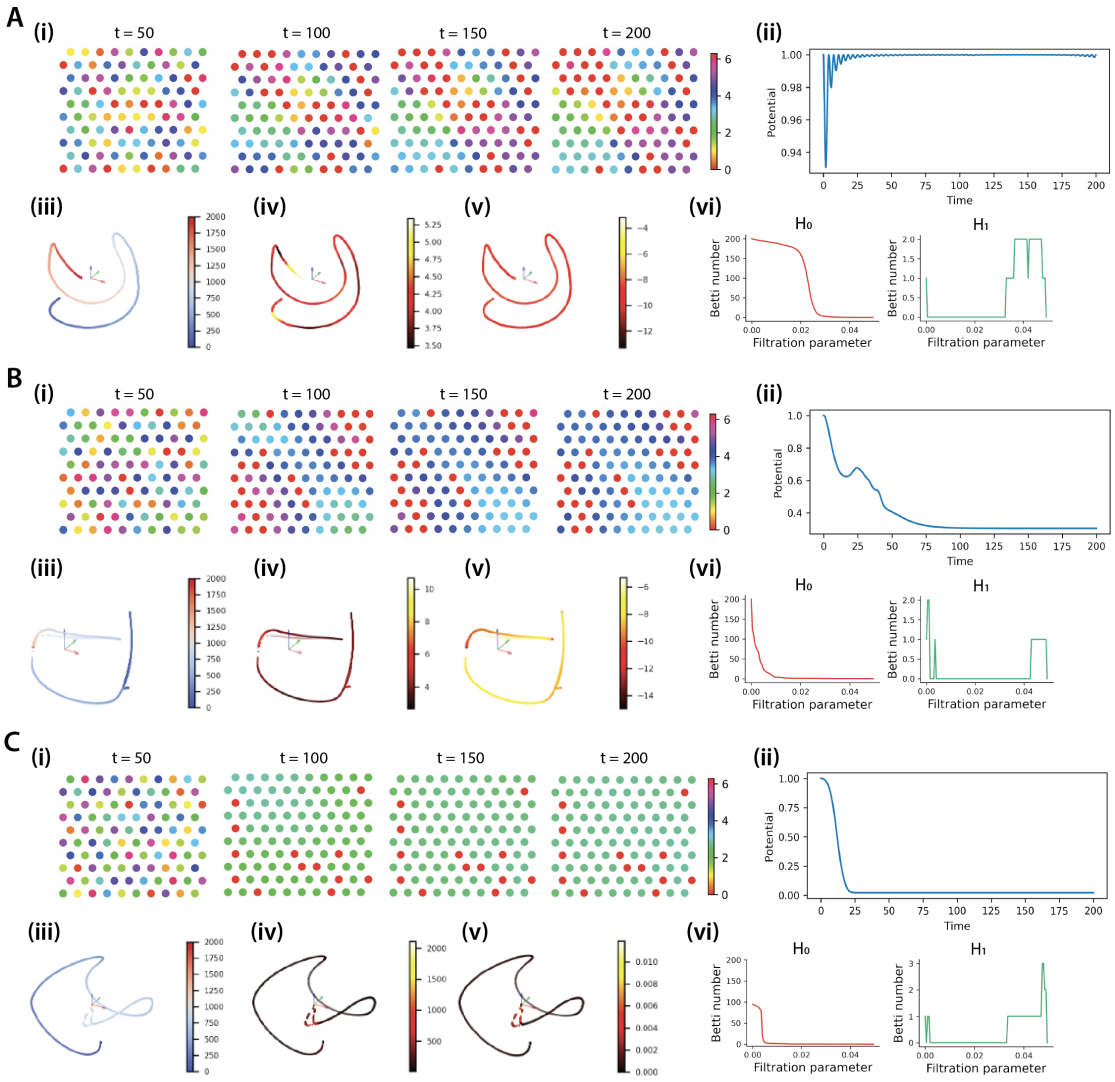

Figure 2: Neuro-GSTH quantifies synchrony and entrainment in simulations of the Kuramoto model.

often visualized as a contour plot, lending to its name, the Contour Realization Of Computed k-dimensional hole Evolution in the Rips complex, or CROCKER (Topaz et al., 2015).

2.2.2 Zig-zag persistent homology

Despite the utility of "ordinary" persistent homology in machine learning and data analysis contexts (Hensel et al., 2021) for a recent survey), the expressivity of such methods is limited by the requirement of data being nested along the filtration, i.e., along the filtering process. In some applications, the use of *zig-zag persistent homology* may thus be preferred. This notion of topological features is permits more general relations between the individual parts of the data. More precisely, given a data set, zig-zag persistent homology can calculate topological features along general sequences of nested data sets of the form

$$\mathrm{K}_1 \leftrightarrow \mathrm{K}_2 \cdots \mathrm{K}_{n-1} \leftrightarrow \mathrm{K}_n, \tag{9}$$

where "$\leftrightarrow$" may either denote subset or superset inclusions, i.e., inclusions of the form "$\hookleftarrow$" or "$\hookrightarrow$." In particular, the types of inclusions are allowed to vary; ordinary persistent homology with sublevel set inclusions is thus a special case of such a zig-zag filtration according to Eqn. 9. The name "zig-zag" denotes the common use case in which both types of inclusions alternate. This can be achieved by inserting the union of data sets, leading to

$$\mathrm{K}_1 \hookrightarrow \mathrm{K}_{1,2} \hookleftarrow \mathrm{K}_2 \hookrightarrow \cdots \hookrightarrow \mathrm{K}_{n-1,n} \hookleftarrow \mathrm{K}_n, \tag{10}$$

where $\mathrm{K}_{i-1,i} := \mathrm{K}_{i-1} \cup \mathrm{K}_i$. It is possible, in complete analogy to ordinary persistent homology, to track topological features along such a zig-zag filtration, making this formulation suitable particularly for time-varying data.

2.2.3 Ablation study

We compare application of the proposed neuro-GSTH method on the three synthetic datasets with approaches that ablate or replace steps in the GSTH method. In particular, we test:

- Applying PHATE directly on the raw input signals to obtain time-trajectories, without the use of the geometric scattering transform.
- Applying PCA on the generated scattering coefficients instead of PHATE.
- Applying t-SNE on the generated scattering coefficients instead of PHATE.
- Applying UMAP on the generated scattering coefficients instead of PHATE.

For the synthetic testcase 1, we aim to compare the approaches for their stability to small perturbations as well as their ability to retrieve signal diffusion dynamics on the graph. As shown in Fig. 3A, after applying neuro-GSTH, time points with perturbed signals overlapped with time points with original signals, showing scattering transform and PHATE are invariant

to small degrees of noise. The smooth trajectory also reflects that the scattering transform and PHATE of neuro-GSTH can effectively capture the signal propagation on the graph. By contrast, directly using PHATE on the raw input signals will result in the condensed timepoints in Fig. 3C, thus failing to retrieve the dynamics. While applying PCA (Fig. 3D) and t-SNE (Fig. 3E) on the generated scattering coefficients can retrieve the dynamics to some extent, Fig. 3D shows a more disrupted trajectory and the trajectory from Fig. 3E overlaps with itself. Similarly, applying UMAP (Fig. 3F) on the generated scattering coefficients also led to overlapping timepoints. All these methods thus failed to reflect the propagation of signals on the graph.

For the second synthetic dataset, we further compare the ability of different approaches to retrieve signal diffusion dynamics on the graph under a more complex scenario. For neuro-GSTH (Fig. 3H) time points from two signal sources formed two branches with their starting points near each other in PHATE coordinates. Thus from one end to the next this is akin to a signal condensing and then diffusing again. As expected, this creates a loop-like structure in the PHATE graph. However, directly applying PHATE on the raw signals (Fig. 3J) results in multiple scattered points separated from the main trajectory, demonstrating that using PHATE only is not able to fully capture and distinguish the signals. Furthermore, although applying PCA on the scattering coefficients (Fig. 3K) generates two separate trajectories, they fail to form the loop-like structure as with using neuro-GSTH. Applying t-SNE (Fig. 3L) and UMAP (Fig. 3M) on the generated scattering coefficients also failed to form loop-like structures.

Finally, for the third synthetic dataset, we aim to simulate the propagation of signals similar to that observed in epithelial cells. This will help us to better understand what types of signals neuro-GSTH can capture. The propagation of signals among cells can reflect the communication within groups of cells (corresponding to the clusters of Ca^{2+} transients among epithelial cells using Ca^{2+}-sensor imaging), while each cell itself can also generate independent signals without communicating with other cells (corresponding to single cells spiking using Ca^{2+} fluorescence imaging). As in Fig. 3O for neuro-GSTH, the time points formed smooth trajectories at first simulating the propagation of signals on the graph, then disrupted trajectories corresponding to random spiking of single cells. In comparison, using PHATE directly on the raw input signals (Fig. 3Q) results in more dense clustering (for the initial stage when the signal is just diffusing on the graph) and using t-SNE on the scattering coefficients generates more scattered clustering, making it hard to identify the inner dynamics (Fig. 3S). Although applying PCA (Fig. 3R) and UMAP (Fig. 3T) on the scattering coefficients can reflect part of the dynamics, they also generate very condensed trajectories for the early stage, when the signal is simply diffusing on the graph.

In addition, we computed the Wasserstein distances between the persistence diagram from our neuro-GSTH method and persistence diagrams from other methods using the three synthetic datasets (Fig. 3G, 3N, 3U). We showed that the persistence diagrams from our GSTH method are different from persistence diagrams produced using other methods ablating different parts of GSTH and visualized the distances with heatmaps (Fig. 3G, 3N, 3U).

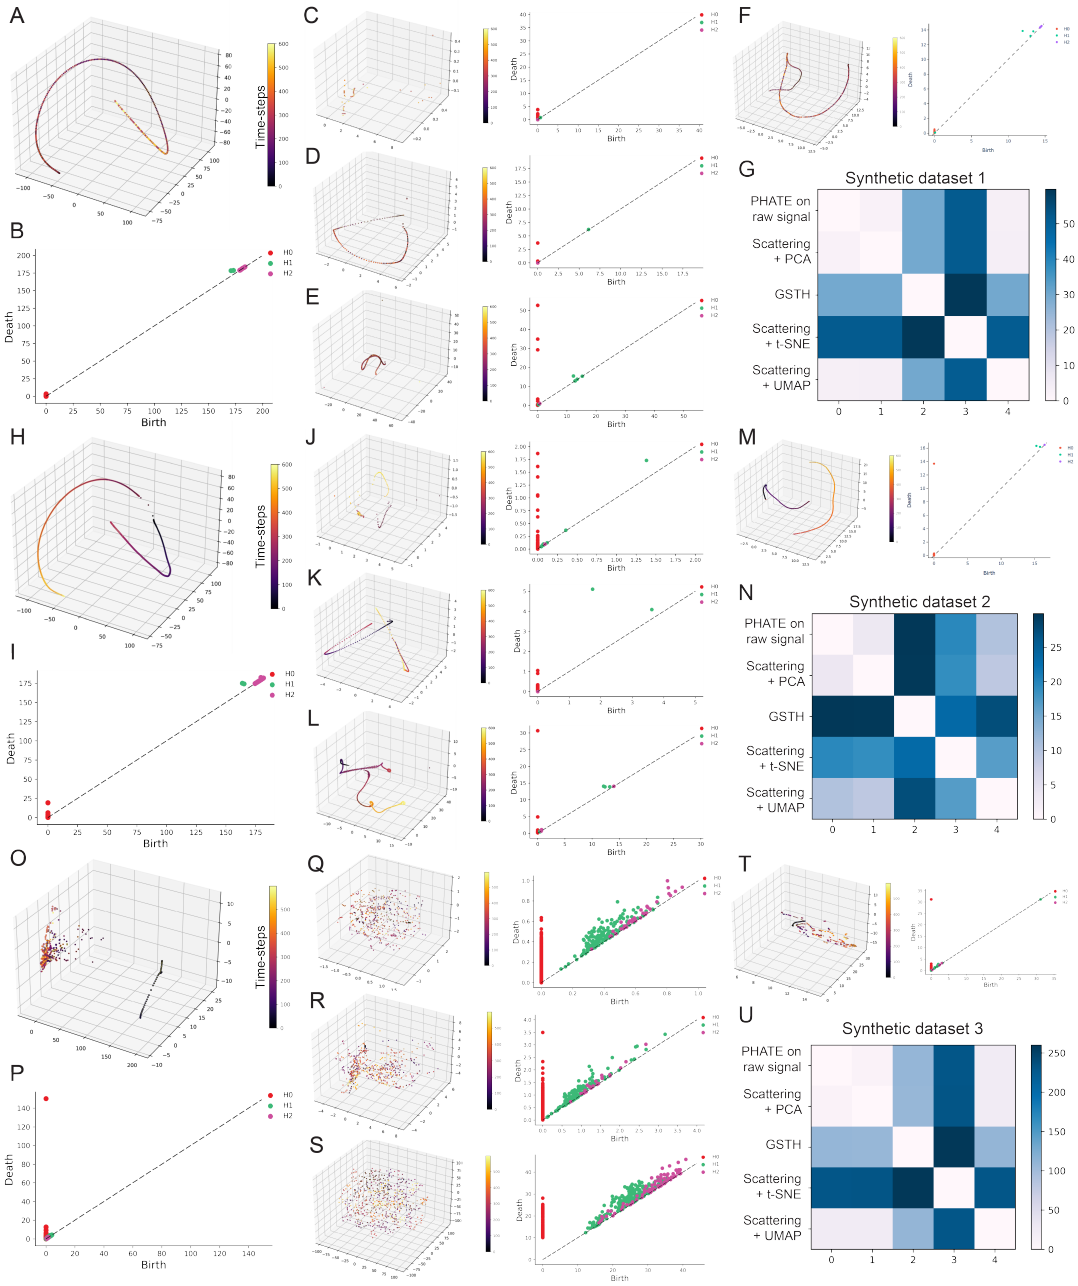

Figure 3: Ablation study

2.3 neuro-GSTH quantifies synchrony in coupled oscillators

We first validated neuro-GSTH by using it to quantify spatiotemporal patterns in the Kuramoto coupled oscillator model. This mathematical model is used to describe collective behavior in oscillating elements and excitable media, such as the firing of neurons in the brain (Kuramoto, 1975; Acebrón et al., 2005; Strogatz, 2000). The model consists of a group of oscillators, each with its own natural frequency, ω, that are coupled together and can influence each other's oscillations. The degree of synchronization among the oscillators is determined by a parameter called

Table 1: Ablation study comparing performance of neuro-GSTH

Model	Kuramoto (MSE)	PVC (Acc.)	fMRI (Acc.)
GSTH	0.0063	0.44	0.64
GST	<u>0.0016</u>	<u>0.62</u>	<u>0.73</u>
GST-RNN	0.0058	0.37	0.59
GS-PCA-H	0.0125	0.44	0.56
GS-tSNE-H	0.0118	0.51	0.48
Graph PH	0.0124	0.53	0.52

the coupling strength, which represents the strength of the interaction between the oscillators.

We simulated a network of 64 coupled oscillators with natural frequencies sampled from a normal distribution, $\omega_i \sim \mathcal{N}(1, \beta)$, and initial phase offsets sampled from a uniform distribution, $\theta_i \sim \mathcal{U}(0, 2\pi)$ for 200 timepoints with $\Delta t = 0.05$. At low coupling strength, the oscillators act independently according to their natural frequency (Fig. 2Ai). We obtain a continuous 1D trajectory using neuro-GSTH, with successive timepoints next to each other in the low-dimensional representation. At intermediate coupling strength the oscillators exhibit entrainment, characterized by decrease in phase variance and increase in phase-locking over time (Fig. 2Aii). The neuro-GSTH trajectory converges over time indicating that most of oscillators eventually are locked in the same state. At high coupling strength, the oscillators synchronize within 20 simulation time-steps (Fig. 2Aiii). Consequently, the neuro-GSTH trajectory converges rapidly, indicating that all oscillators share the same state. The shape of the neuro-GSTH trajectories is quantified using Betti curves, with clear differences in dimension 0 persistent homology showing the increasing convergence of the tPHATE embeddings with increase in synchrony and entrainment. We are able to recover of the coupling strength parameter of the simulations from neuro-GSTH embeddings and topological descriptors as shown in Table 1.

2.4 neuro-GSTH reconstructs sinusoidal patterns shown to mice from visual cortex recordings

Using publicly available two-photon microscopy data (Garner, 2014a), we showed that neuro-GSTH can recover sinusoidal grating patterns presented to mice from calcium imaging of layer 4 regions of the prefrontal visual cortex. Sinusoidal grating patterns are a type of visual stimulus consisting of parallel stripes that vary in brightness according to a sinusoidal function. These patterns are commonly used in vision research because they are simple, well-defined stimuli that can evoke robust responses in the visual cortex.

Moving striped patterns on a computer screen with varying spatial and temporal frequencies and orientations were shown to mice, fitted with head-plates and a cranial window. Each pattern was shown for 3 seconds with a 2-second break in between, and this was repeated 8 times for each pattern. The monitor was positioned in front of the mouse's eye and the images were

collected from the layer 4 region in the visual cortex, at a rate of approximately 30 images per second. We treat the task of recovering the sinusoidal visual stimuli presented to mice as a classification task. Specifically, we bin the different spatial frequencies, temporal frequencies, and orientations of the sinusoidal patterns into 5 distinct categories and train a a 5-layer feed-forward neural network to predict the correct category for each stimulus type from neuro-GSTH trajectories.

Our classifier significantly outperforms a random forest classifier in predicting the spatial frequency of the sinusoidal patterns, achieving a classification accuracy of 0.78 ± 0.01 compared to the random forest's accuracy of 0.64 ± 0.01. In predicting the orientation of sinusoidal patterns, neuro-GSTH also outperforms the random forest, though the performance gap is smaller, with neuro-GSTH achieving 0.72 ± 0.01 accuracy versus 0.63 ± 0.02 for the random forest. Finally, neither our model nor the random forest can accurately classify the temporal frequency of the sinusoidal patterns, suggesting that this information may not be represented in the layer 4 region of the visual cortex that was imaged in this study. One possible explanation is that temporal frequency information is not represented in the layer 4 region of the visual cortex which was imaged in this study. Importantly, neuro-GSTH outperformed other methods of quantifying the signaling dynamics in all cases.

3 Conclusion

In conclusion, this study demonstrates the remarkable potential of combining geometric scattering and persistent homology for the quantification and analysis of spatiotemporal signaling in various biological systems. Through the successful prediction of synchrony in coupled oscillators and the recovery of sinusoidal visual stimuli presented to mice from calcium imaging of the visual cortex, neuro-GSTH has showcased its capability to capture intricate patterns and reveal hidden information within neural processes.

Overall, the integration of geometric scattering and persistent homology within a single framework offers a powerful and versatile tool for studying spatiotemporal patterns. The demonstrated success across various biological systems and applications underscores the potential of this framework in advancing our understanding of cellular communication, coordination, and regulation. By unlocking new insights into the complex dynamics of neural communication, this innovative approach holds great promise for advancing our understanding of diverse biological systems and has the potential to transform the fields of cognitive and systems neuroscience, psychiatry, and psychology, among others.

4 Methods

4.1 Mathematical background

4.1.1 Cellular Graphs and Graph Signals

We represent the imaged tissue as a graph $G = \{V, E\}$, consisting of nodes $v_i \in V$ and edges $(v_j, v_k) \in E$, where each node v_i represents a cell and a pair of nodes v_j and v_k is connected with an edge based on a predefined criterion. For epithelial cells, we connect nodes that are spatially adjacent (within 2 μm of each other), as the flow of signals is thought to be between spatially proximal cells. On the other hand, neurons can have long processes that are often hard to image, and therefore we use correlation between neurons' Ca^{2+} signals to connect the neuronal graph. Finally, the connectivity of graph G can be described by its adjacency matrix $\mathbf{A}$, where $\mathbf{A}_{ij} = 1$ if v_i and v_j are connected and 0 otherwise. The degree of each vertex is defined as a diagonal matrix $\mathbf{D}$, where $\mathbf{D}_{ii} = \sum_{j=1}^{i} \mathbf{A}_{ij}$.

Graph signals can associate with each node or edge in a graph. In the Ca^{2+} signaling data, the signals associated with cell v_i is the normalized Ca^{2+} fluorescence intensity at each timestep t. Since every cell has a related Ca^{2+} signal, this signal $X(v_i, t)$ is defined over the whole graph for timestep t.

4.1.2 Diffusion Geometry

A useful assumption in representation learning is that high dimensional data originates from an intrinsic low dimensional manifold that is mapped via nonlinear functions to observable high dimensional measurements; this is commonly referred to as the manifold assumption. Formally, let $\mathcal{M}^d$ be a hidden d-dimensional manifold that is only observable via a collection of $n \gg d$ nonlinear functions $f_1, \ldots, f_n : \mathcal{M}^d \to \mathbb{R}$ that enable its immersion in a high dimensional ambient space as $F(\mathcal{M}^d) = \{\mathbf{f}(z) = (f_1(z), \ldots, f_n(z))^T : z \in \mathcal{M}^d\} \subseteq \mathbb{R}^n$ from which data is collected. Conversely, given a dataset $X = \{x_1, \ldots, x_N\} \subset \mathbb{R}^n$ of high dimensional observations, manifold learning methods assume data points originate from a sampling $Z = \{z_i\}_{i=1}^N \in \mathcal{M}^d$ of the underlying manifold via $x_i = \mathbf{f}(z_i)$, $i = 1, \ldots, n$, and aim to learn a low dimensional intrinsic representation that approximates the manifold geometry of $\mathcal{M}^d$.

To learn a manifold geometry from collected data, scientists often use the diffusion maps construction of (Coifman and Maggioni, 2006) that uses diffusion coordinates to provide a natural global coordinate system derived from eigenfunctions of the heat kernel, or equivalently the Laplace-Beltrami operator, over manifold geometries. This construction starts by considering local similarities defined via a kernel $\mathcal{K}(x, y)$, $x, y \in F(\mathcal{M}^d)$, that captures local neighborhoods in the data. We note that a popular choice for $\mathcal{K}$ is the Gaussian kernel $\exp(-\|x-y\|^2/\sigma)$, where $\sigma > 0$ is interpreted as a user-configurable neighborhood size. However, such neighborhoods encode sampling density information together with local geometric information. To construct a

diffusion geometry that is robust to sampling density variations, we use an anisotropic kernel:

$$\mathcal{K}(x,y) = \frac{\mathcal{G}(x,y)}{\|\mathcal{G}(x,\cdot)\|_1^{\alpha}\|\mathcal{G}(y,\cdot)\|_1^{\alpha}}, \text{ where } \quad \mathcal{G}(x,y) = e^{-\frac{\|x-y\|^2}{\sigma}}, \tag{11}$$

as proposed in Coifman and Maggioni (2006), where $0 \leq \alpha \leq 1$ controls the separation of geometry from density, with $\alpha = 0$ yielding the classic Gaussian kernel, and $\alpha = 1$ completely removing density and providing a geometric equivalent to uniform sampling of the underlying manifold. Next, the similarities encoded by $\mathcal{K}$ are normalized to define transition probabilities $p(x,y) = \frac{\mathcal{K}(x,y)}{\|\mathcal{K}(x,\cdot)\|_1}$ that are organized in an $N \times N$ row stochastic matrix:

$$\mathbf{P}_{ij} = p(x_i, x_j) \tag{12}$$

that describes a Markovian diffusion process over the intrinsic geometry of the data. Finally, a diffusion map is defined by taking the eigenvalues $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ and (corresponding) eigenvectors $\{\phi_j\}_{j=1}^N$ of $\mathbf{P}$, and mapping each data point $x_i \in X$ to an N dimensional vector $\Phi_t(x_i) = [\lambda_1^t \phi_1(x_i), \ldots, \lambda_N^t \phi_N(x_i)]^T$, where t represents a diffusion-time (i.e., number of transitions considered in the diffusion process). In general, as t increases, most of the eigenvalues λ_j^t, $j = 1, \ldots, N$, become negligible, and thus truncated diffusion map coordinates can be used for dimensionality reduction (Coifman and Maggioni, 2006).

4.1.3 PHATE

PHATE is a dimensionality reduction method that captures both local and global nonlinear structure through constructing a diffusion geometry (Moon et al., 2019). It computes the diffusion operator as in Eqn. 12. However, rather than eigendecomposing this operator to find new coordinates, PHATE creates a new distance matrix from $\mathbf{P}$ by defining an M-divergence between datapoints, called *potential distance* as $\mathbf{ID}_{i,j} = \| \log \mathbf{P}_{i,:}^t - \log \mathbf{P}_{j,:}^t \|_2$ between corresponding t-step diffusion probability distributions of the two points.

The advantage of this step is that the information theoretic distance between probabilities emphasizes differences in lower probabilities (corresponding to distant points) as well as high probabilities (corresponding to neighbors), and therefore globally contextualizes the point. The resulting information distance matrix $\mathbf{ID}$ is finally embedded into a low dimensional (2D or 3D) space by metric multidimensional scaling (MDS), and makes it possible to visualize intrinsic geometric information from data.

4.1.4 Persistent homology and topological data analysis

Topological data analysis (TDA) refers to techniques for understanding complex datasets by their topological features, i.e., their connectivity (Carlsson, 2009). Here we focus on the topological features of a data graph where the simplest set of topological features are given by the number of connected components b_0 and the number of cycles b_1, respectively. Such counts, also known as the Betti numbers, are coarse graph descriptors that are invariant under graph isomorphisms. Their expressivity is increased by considering a function $f\colon V \times V \to \mathbb{R}$ on the

vertices of a graph, $G = (V, E)$, with vertex set V and edge set E. Since V has finite cardinality, so does the image im$f$, i.e., im$f = \{w_1, w_2, \ldots, w_n\}$.

Without loss of generality, we assume that $w_1 \leq \cdots \leq w_n$. We write G_i for the subgraph induced by filtering according to w_i, such that the edges satisfy $E_i := \{(u, v) \in E \mid \|u - v\|_2^2 \leq w_i\}$. The subgraphs G_i satisfy a nesting property, as $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_n$. When analyzing a point cloud, the vertices of each G_i arise from spatial coordinates for the data and w_i constitutes a distance threshold between points, such that G_n is a fully-connected graph, containing all the vertices from V. This is commonly known as the Vietoris-Rips (VR) filtration.

It is then possible to calculate topological features alongside this *filtration* of graphs, tracking their appearance and disappearance as the graph grows. If a topological feature is created in G_i, but destroyed in G_j (it might be destroyed because two connected components merge, for instance), we represent this by storing the point (w_i, w_j) in the *persistence diagram* $\mathcal{D}_f$ associated to G. Another simple descriptor is given by the Betti curve of dimension d of a diagram $\mathcal{D}$, which refers to the sequence of Betti numbers of dimension d in $\mathcal{D}$, evaluated for each threshold w_i.

4.2 Kuramoto model of weakly coupled oscillators

The Kuramoto model (Pikovsky and Rosenblum, 2015) describes a weakly coupled system of oscillators, with dynamics governed by:

$$\dot{\phi_k} = \omega_k + \frac{\epsilon}{N} \sum_{j=1}^{N} \sin(\phi_j - \phi_k + \alpha_k) \tag{13}$$

where ω_k are the natural frequencies and α_k are the phase shift parameters.

By using the complex mean-field, $Z = Re^{i\theta} = {}^1\!/_N \sum_j e^{i\phi_j}$, the system can be reduced to:

$$\dot{\phi_k} = \omega_k + \epsilon R \sin(\theta - \phi_k + \alpha_k) \tag{14}$$

The order parameter R describes the degree of synchrony, $R = 1$ denotes full synchronization and at $R = 0$ each oscillator acts independently.

4.3 Experimental setup for prefrontal visual cortex dataset

Image data generated in Garner (2014b) were downloaded from the Collaborative Research in Computational Neuroscience (CRCNS) data sharing website. The image data were acquired using a two-photon microscope and images were collected at approximately 30 Hz with 512 lines per frame. Adult mice were fitted with head plates and a cranial window and allowed to run on a freely rotating disc while visual stimuli were presented to the eye contralateral to the imaged region (layer 4 cells of the primary visual cortex). Drifting sinusoidal gratings were shown in random order for all combinations of 5 spatial frequencies, 5 temporal frequencies, and 8 orientations (0 to 315 degrees in 45 degree steps). These stimuli were presented 8 times each for 3 seconds.

References

Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F., and Spigler, R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1):137–185. Publisher: American Physical Society.

Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308.

Coifman, R. R. and Maggioni, M. (2006). Diffusion wavelets. *Applied and Computational Harmonic Analysis*, 21(1):53–94.

Garner, A. (2014a). In vivo calcium imaging of layer 4 cells in the mouse using sinusoidal grating stimuli.

Garner, A. (2014b). In vivo calcium imaging of layer 4 cells in the mouse using sinusoidal grating stimuli.

Hensel, F., Moor, M., and Rieck, B. (2021). A survey of topological machine learning methods. *Frontiers in Artificial Intelligence*, 4.

Kuramoto, Y. (1975). International symposium on mathematical problems in theoretical physics. *Lecture notes in Physics*, 30:420.

Moon, K. R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D. B., Chen, W. S., Yim, K., van den Elzen, A., Hirn, M. J., Coifman, R. R., Ivanova, N. B., Wolf, G., and Krishnaswamy, S. (2019). Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology*, 37(12):1482–1492.

Moore, J. L., Gao, F., Matte-Martone, C., Du, S., Lathrop, E., Ganesan, S., Shao, L., Bhaskar, D., Cox, A., Hendry, C., Rieck, B., Krishnaswamy, S., and Greco, V. (2021). Tissue-wide coordination of calcium signaling regulates the epithelial stem cell pool during homeostasis. Pages: 2021.10.12.464066 Section: New Results.

Pikovsky, A. and Rosenblum, M. (2015). Dynamics of globally coupled oscillators: Progress and perspectives. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(9):097616. Publisher: American Institute of Physics.

Strogatz, S. H. (2000). From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1):1–20.

Topaz, C. M., Ziegelmeier, L., and Halverson, T. (2015). Topological Data Analysis of Biological Aggregation Models. *PLOS ONE*, 10(5):e0126383. Publisher: Public Library of Science.

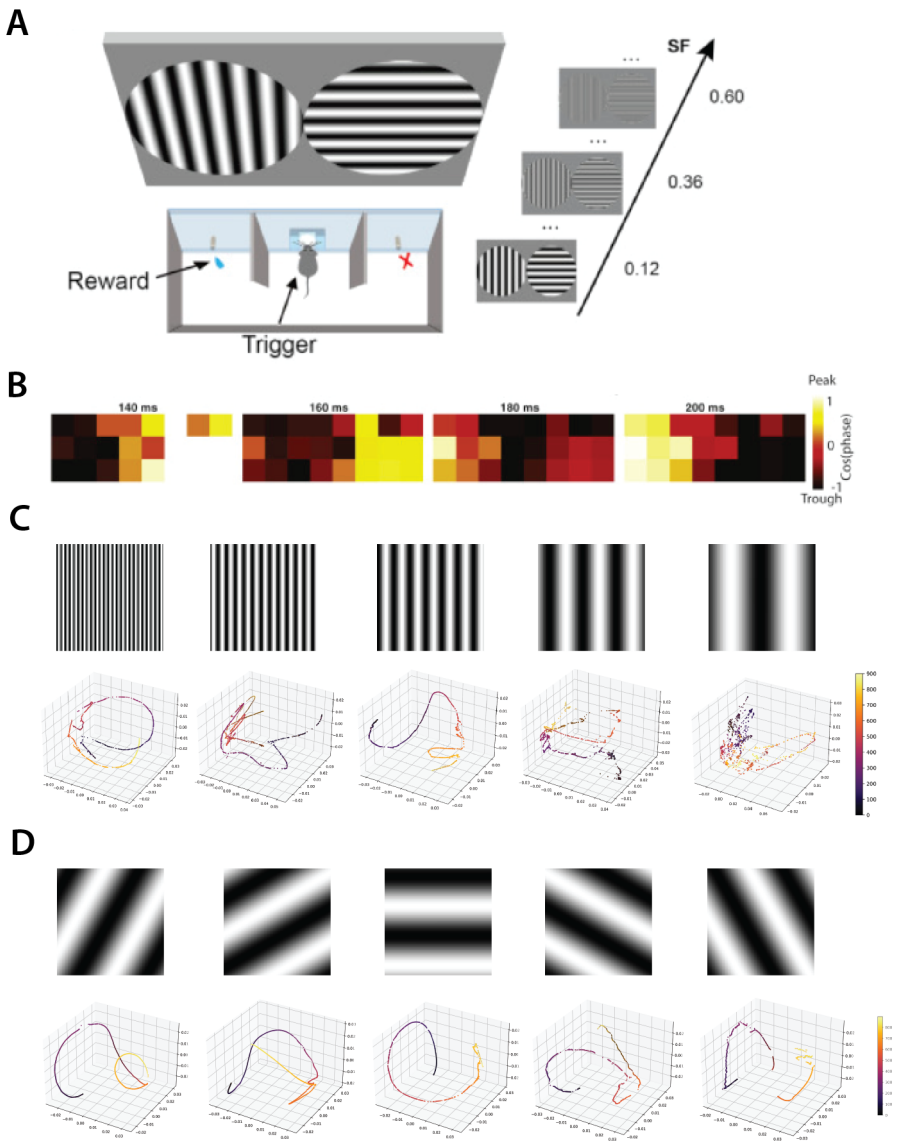

Figure 4: Signatures of neural activity in layer 4 of the primary visual cortex in response to sinusoidal visual stimuli at varying spatial frequency and orientations.